arXiv:1404.3131v4 [cs.DB] 22 Jul 2014

The Possibility Problem for Probabilistic XML (Extended Version)

Antoine Amarilli
Télécom ParisTech; Institut Mines-Télécom; CNRS LTCI

Abstract. We consider the *possibility problem* of determining if a document is a possible world of a probabilistic document, in the setting of probabilistic XML. This basic question is a special case of query answering or tree automata evaluation, but it has specific practical uses, such as checking whether an user-provided probabilistic document outcome is possible or sufficiently likely.
In this paper, we study the complexity of the possibility problem for probabilistic XML models of varying expressiveness. We show that the decision problem is often tractable in the absence of long-distance dependencies, but that its computation variant is intractable on unordered documents. We also introduce an *explicit matches* variant to generalize practical situations where node labels are unambiguous; this ensures tractability of the possibility problem, even under long-distance dependencies, provided event conjunctions are disallowed. Our results entirely classify the tractability boundary over all considered problem variants.

1 Introduction

Probabilistic representations are a way to represent incomplete knowledge through a concise description of a large set of possible worlds annotated with their probability. Such models can then be used, e.g., to run a query efficiently over all possible worlds and determine the overall probability that the query holds. Probabilistic representations have been successfully used both for the relational model [20] and for XML documents [16].

Many problems, such as query answering [15], have been studied over such representations; however, to our knowledge, the *possibility problem* (POSS) has not been specifically studied: given a probabilistic document D and a deterministic document W, decide if W is a possible world of D, and optionally compute its probability according to D. This can be asked both of relational and XML probabilistic representations, but we focus on XML documents because they pose many challenges: they are hierarchical so some probabilistic choices appear dependent[1]; documents may be ordered; bag semantics must be used to count multiple sibling nodes with the same label. In addition, in the XML setting, the POSS problem is a natural question that arises in practical scenarios.

As a first example, when using probabilistic XML to represent a set D of possible versions [6] of an XML document, one may want to determine if a version W, obtained from a user or from an external source, is one of the known possible versions represented as a probabilistic XML document D. For instance, assume that a probabilistic

[1] In fact, we will see that our hardness results always hold even for shallow documents.

XML version control system asks a user to resolve a conflict [5], whose uncertain set of possible outcomes is represented by D. When the user provides a candidate merge W, the system must then check if the document W is indeed a possible way to solve the conflict. This may be hard to determine, because D may, in general, have many ways to generate W, through a possibly intractable number of different valuations of its uncertainty events.

As a second practical example, assume that a user is studying an uncertain document D that provides a representation of possible versions of an XML tree, using probabilistic XML to represent possible conflicting choices and their probability. The user notices that choosing a certain combination of decisions yields a certain deterministic document W, and asks whether the same document could have been obtained by making different choices. Indeed, maybe W is considered improbable under D following this particular valuation, but is likely overall because the same document can be obtained through many different ways. What is the probability, over all valuations, of the user's chosen outcome W according to D?

On the face of it, POSS seems related to query evaluation: we wish to evaluate on D a query q_W which is, informally, "is the input document exactly W"? However, there are three reasons why query evaluation cannot give good complexity bounds for POSS. First, because q_W depends on the possibly large W, we are not performing query answering for a *fixed* query, so we can only use the unfavorable *combined complexity* bounds where both the input document D and the query q_W are part of the input. Second, because we want to obtain *exactly* W, the match of q_W should never map two query variables to the same node of D, so the query language must allow inequalities on node identifiers. Third, once again because we require an exact match, we need to assert the *absence* of the nodes which are not in W, so we need negation in the language. To our knowledge, then, the only upper bound for POSS from query answering is the combined complexity bound for the (expressive) *monadic second-order logic over trees* whose evaluation on deterministic (not even probabilistic) XML trees is already PSPACE-hard [18].

A second related approach is that of tree automata on probabilistic XML documents. Indeed, we can encode the possible world W to a deterministic tree automaton A_W and compute the probability that A_W accepts the probabilistic document D. The decision and computation variants of POSS under local uncertainty models are thus special cases of the "relevancy" and "p-acceptance" problems of [9]. However, their work only considers *ordered* trees, and an unordered W cannot easily be translated to their deterministic tree automata, because of possible label ambiguity: we cannot impose an arbitrary order on D and W, as this also chooses how nodes must be disambiguated. In fact, we will show that POSS is hard in some settings that are tractable for ordered documents.

This paper specifically focuses on the POSS problem to study the precise complexity of its different formulations. Our probabilistic XML representation is the PrXML model of [16], noting that some results are known for the POSS problem (called the "membership problem") in the incomparable and substantially different "open-world" incomplete XML model of [8] (whose documents have an infinite set of possible worlds, instead of a possibly exponential but finite set as in PrXML).

We start by defining the required preliminaries in Section 2 and the different variants of POSS in Section 3, establishing its overall NP-completeness and reviewing the

results of [9]. We then study local uncertainty models in Section 4 and show that the absence of order impacts tractability, with a different picture for the decision and computation variants of POSS. Last, in Section 5, we show that POSS can be made tractable under long-distance event correlations, by disallowing event conjunctions and imposing an "explicit matches" condition which generalizes, e.g., unique node labels. We then conclude in Section 6.

This paper is the complete version (including proofs) of work initially submitted as an extended abstract (without proofs) at the AMW 2014 workshop [3] and subsequently submitted (with proofs) at the BDA 2014 conference (no formal proceedings). This version integrates the feedback from both rounds of reviews.

2 Preliminaries

We start by formally defining XML documents and probability distributions over them:

Definition 1. *An* unordered XML document *is an unordered tree whose nodes carry a label from a set Λ of labels.* Ordered *XML documents are defined in the same way but with ordered trees, that is, there is a total order over the children of every node.*

A probability distribution *is a function $\mathcal{P}$ mapping every XML document x from a finite set* $\mathrm{supp}(\mathcal{P})$ *to a rational number $\mathcal{P}(x)$, its* probability *according to $\mathcal{P}$, with the condition that $\sum_{D \in \mathrm{supp}(\mathcal{P})} \mathcal{P}(D) = 1$. For any $x \notin \mathrm{supp}(\mathcal{P})$ we write $\mathcal{P}(x) = 0$.*

As it is unwieldy to manipulate explicit probability distributions over large sets of documents, we use the language of probabilistic XML [16] to write extended XML documents (with so-called *probabilistic nodes*) and give them a semantics which is a (possibly exponentially larger) probability distribution over XML documents. Intuitively, probabilistic XML documents are XML documents with specific *probabilistic* nodes describing possible choices in the document; their semantics is the set of XML documents that can be obtained under those choices.

Definition 2. *A* PrXML probabilistic XML document *D is an XML document over $\Lambda \sqcup \{\mathsf{det}, \mathsf{ind}, \mathsf{mux}, \mathsf{cie}, \mathsf{fie}\}$. The nodes of D with labels from Λ are called* regular *nodes, by opposition to* probabilistic *nodes. The probabilistic labels respectively stand for: determininistic, independent, mutually exclusive, conjunction of independent events, formula of independent events. For any subset $\mathcal{L} \subseteq \{\mathsf{det}, \mathsf{ind}, \mathsf{mux}, \mathsf{cie}, \mathsf{fie}\}$, we call $\mathsf{PrXML}^{\mathcal{L}}$ the language of probabilistic XML documents containing only nodes with labels in $\Lambda \sqcup \mathcal{L}$.*

We require that the root of a PrXML *document D is a regular node, that every edge from a* mux *or* ind *node to a child node is labeled with some rational number*[2] *$0 < x < 1$ (the sum of the labels of the children of every* mux *node being ≤ 1), and that every edge from a* cie *(resp.* fie*) node to a child node is labeled with a conjunction (resp. a Boolean formula) of* events *from a set E of events (and their negations), with D providing a mapping $\pi : E \to [0,1]$ attributing a rational probability to every event.*

[2] The non-standard constraint $x < 1$ means that ind does not subsume det (see Thm. 3 and 4 for examples where this distinction matters).

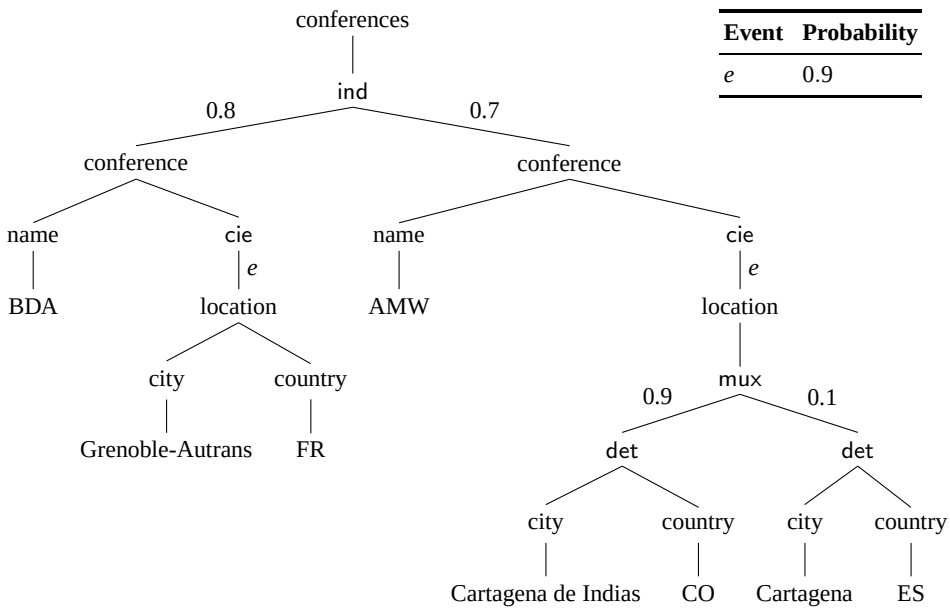

Event	Probability
e	0.9

Fig. 1: Example $\mathsf{PrXML}^{\mathsf{mux,ind,det,cie}}$ document; the provided table is the mapping π that attributes probabilities to probabilistic events

The semantics of a PrXML document D is the probability distribution over XML documents defined by the following sampling process (see [16] for more details):

Definition 3. *A deterministic XML document W is obtained from a* PrXML *document D as follows. First, choose a valuation $\nu : E \to \{\mathfrak{t},\mathfrak{f}\}$ of the events from E, with probability $\prod_{e \text{ s.t.} \nu(e)=\mathfrak{t}} \pi(e) \times \prod_{e \text{ s.t.} \nu(e)=\mathfrak{f}} (1-\pi(e))$. Evaluate* cie *and* fie *nodes by keeping only the child edges whose Boolean formula is true under ν. Evaluate* ind *nodes by choosing to keep or delete every child edge according to the probability indicated on its edge label. Evaluate* mux *nodes by removing all of their children edges, except one chosen according to its probability (possibly keep none if the probabilities sum up to less than* 1*). Finally, evaluate* det *nodes by replacing them by the collection of their children.*

All probabilistic choices are performed independently, so the overall probability of an outcome is the product of the probabilities at each step. Whenever an edge is removed, all of the descendant nodes and edges are removed. The probability of a document W according to D, written $D(W)$, is the total probability of all outcomes[3] *leading to W.*

We say that mux, ind and det are *local* in the sense that they describe a probabilistic choice that *takes place* at this point of the document, independently from other choices (except for the fact that discarding a subtree makes irrelevant all local probabilistic

[3] Note that in general there may be multiple outcomes that lead to the same document W.

Problem		Complexity	
POSS	$\top$ fie	NP	(Prop. 1)
#POSS	$\top$ fie	$FP^{\#P}$	(Prop. 1)
#POSS	$<$ mux, ind, det	PTIME	(Thm. 1)
#POSS	$\not<$ ind or mux	#P-hard	(Thm. 2)
POSS	$\top$ ind or mux	PTIME	(Thm. 3)
POSS	$\not<$ 2 of mux, ind, det	NP-hard	(Thm. 4)
#EPOSS	$\top$ mux, ind, det	PTIME	(Thm. 5)
EPOSS	$\bot$ cie	NP-hard	(Thm. 6)
POSS	$\bot$ mie	NP-hard	(Thm. 7)
#EPOSS	$\top$ mie	PTIME	(Thm. 8)

Table 1: Summary of results

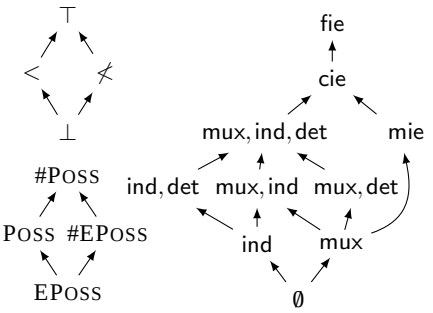

Fig. 2: Variants and PTIME reductions

choices in that subtree. By contrast, we say that cie and fie are *long-distance* in the sense that a valuation is chosen *globally* for the probabilistic events and the cie and fie nodes are then evaluated according to that choice: this may induce *correlations* between arbitrary portions of the document, because the same event can be reused multiple times at different positions in the document.

Example 1. Consider the example probabilistic XML document D in Figure 1. Its possible worlds are obtained as follows. First, draw a valuation for the (only) event e, which may be $\mathfrak{f}$ (with probability 0.1) or $\mathfrak{t}$ (with probability 0.9). Then, decide whether to keep or discard the first "conference" subtree, with probability 0.8, and decide whether to keep or discard the second such subtree, with probability 0.7. Remove the cie nodes and keep or discard their children depending on whether the chosen valuation for e is $\mathfrak{t}$ or $\mathfrak{f}$ respectively. Decide whether to keep the first or second child of the mux node, and replace the corresponding det node by its children. All probabilistic choices are made independently.

Observe how the choice on mux is irrelevant if the corresponding subtree was discarded by the parent ind node or by the cie node, and notice the use of det nodes to switch between sets of nodes using a mux node. Note that the use of cie nodes introduces a *correlation* in the sense that the first "location" node is present if and only if the second is also present.

Of course, the expressiveness and compactness of PrXML frameworks depend on which probabilistic nodes are allowed: we say that $\mathsf{PrXML}^{\mathcal{C}}$ is *more general* than $\mathsf{PrXML}^{\mathcal{D}}$ if there is a polynomial time algorithm to rewrite any $\mathsf{PrXML}^{\mathcal{D}}$ document to a $\mathsf{PrXML}^{\mathcal{C}}$ document representing the same probability distribution. Fig. 2 (adapted from [14]) represents this hierarchy on the PrXML classes that we consider.

3 Problem and general bounds

We now define the POSS problem formally, in its decision and computation variants.

Definition 4. *Given a class* $\mathsf{PrXML}^{\mathcal{C}}$*, the* possibility problem *for* unordered *documents* $\text{POSS}^{\mathcal{C}}_{\not<}$ *is to determine, given as input an unordered* $\mathsf{PrXML}^{\mathcal{C}}$ *document* D *and an unordered XML document* W*, whether* W *is a possible world of* D*, namely,* $D(W) > 0$.

The possibility problem *for* ordered *documents* $\textsc{Poss}^{\mathcal{C}}_{<}$ *is the same problem except that both D and W are ordered. For $o \in \{\not<, <\}$, the* $\#\textsc{Poss}^{\mathcal{C}}_{o}$ *problem is to compute the probability $D(W)$ of W according to D. Observe that* $\#\textsc{Poss}^{\mathcal{C}}_{o}$ *is a* computation problem *rather than a decision problem, namely, it computes an output value based on the provided input (here, a probability value) instead of merely deciding whether to accept or reject.*

For brevity, we write $\textsc{Poss}^{\mathcal{C}}_{\bot}$ and $\textsc{Poss}^{\mathcal{C}}_{\top}$ when describing lower or upper complexity bounds that apply to both $\textsc{Poss}^{\mathcal{C}}_{<}$ and $\textsc{Poss}^{\mathcal{C}}_{\not<}$.

We start by giving straightforward bounds on the most general problem variants:

Proposition 1. $\textsc{Poss}^{\mathsf{fie}}_{\top}$ *is in NP and* $\#\textsc{Poss}^{\mathsf{fie}}_{\top}$ *is in* $\mathrm{FP}^{\#\mathrm{P}}$.

Proof. We first show the NP-membership of $\textsc{Poss}^{\mathsf{fie}}_{\top}$.

Let us first consider $\textsc{Poss}^{\mathsf{fie}}_{<}$. Consider the input (D, W). Guess a valuation of the probabilistic events of D. The size of the guess is linear in $|D|$. Now, check that the guess is suitable, namely, that the deterministic document D' obtained from D under this valuation is exactly W: as both D' and W are totally ordered trees, this can be checked straightforwardly in linear time from a simultaneous traversal of D' and W. Hence, $\textsc{Poss}^{\mathsf{fie}}_{<}$ is in NP.

Let us now consider $\textsc{Poss}^{\mathsf{fie}}_{\not<}$. The proof idea is the same, except that checking that D' and W are equal is not as obvious, because those trees are not ordered; however, this check can be performed in PTIME by a dynamic bottom-up algorithm similar to that of the proof of Thm. 3, so that the result still holds.

We now show the $\mathrm{FP}^{\#\mathrm{P}}$-membership of $\#\textsc{Poss}^{\mathsf{fie}}_{\top}$.

We first preprocess all the event probabilities in the probabilistic document D so that all numbers are represented with the same denominator. This can be done in polynomial time by a least common multiple computation and product operations. We then read off the common denominator, d. We can compute d^k in polynomial time, where k is the number of events.

We then use Lemma 5.2 of [1] to argue that it is possible, in #P, to compute the unnormalized probability of W, that is, the probability of W in D without dividing by d^k. To do so, the generating PTIME Turing machine T enumerates all possible valuations, and the function g returns 0 if the outcome does not yield the desired document W (which can be decided in PTIME by the above proof for the decision problem), and otherwise returns the unnormalized probability of the outcome, that is, the product of the numerators of the involved probabilities. (The denominator, which would be d^k, is ignored for now.) Hence, by application of this lemma, $\#\textsc{Poss}^{\mathsf{fie}}_{\top}$ is in $\mathrm{FP}^{\#\mathrm{P}}$, because all that remains is to divide the result of this #P computation by d^k to obtain the final probability. □

Proposition 2. $\textsc{Poss}^{\mathsf{cie}}_{\bot}$ *is NP-complete, even when D has height* 3.

Proof. From Prop. 1 it suffices to show hardness. We show a reduction from the NP-hard Boolean satisfiability problem to justify that $\textsc{Poss}^{\mathsf{cie}}_{\bot}$ is NP-hard.

Consider a formula F formed of a conjunction of disjunctive clauses $(C_i)_{1 \leq i \leq n}$, with clause C_i containing the literals $(l^i_j)_{1 \leq j \leq n_i}$, each literal being a positive or negative occurrence of some variable from a finite set of variables $V = \{x_1, \ldots, x_m\}$.

Consider a set of m Boolean events E with a mapping ϕ associating $x_i \in V$ to $e_i \in E$ (and $\neg x_i$ to $\neg e_i$) for all $1 \leq i \leq m$. Consider W the document with only one root labeled $\top$, and the $\mathsf{PrXML}^{\mathsf{cie}}$ document D, with events E (and probability $1/2$ for each outcome), with one root labeled $\top$ and one cie child with n children labeled $\bot$, the edge of the i-th child being labeled with $C'_i = \neg\phi(l^i_1) \wedge \cdots \wedge \neg\phi(l^i_{n_i})$. Given the shape of W, clearly the algorithm's choice to consider D and W as either ordered or unordered trees is irrelevant, so that this works as a reduction either to $\textsc{Poss}^{\mathsf{cie}}_{\not<}$ or to $\textsc{Poss}^{\mathsf{cie}}_{<}$.

Now, W is a possible world of D if and only if there is a valuation of the events of E such that $\bigwedge_i \neg C'_i$ holds, or, equivalently by De Morgan's law, such that $\bigwedge_i \bigvee_j \phi(l^i_j)$ holds, hence (D, W) is a positive instance of $\textsc{Poss}^{\mathsf{cie}}_{\bot}$ if and only if F is satisfiable. Hence, $\textsc{Poss}^{\mathsf{cie}}_{\not<}$ is NP-hard. □

Local models on ordered documents are known to be tractable using tree automata:

Theorem 1 ([9]). $\#\textsc{Poss}^{\mathsf{mux,ind,det}}_{<}$ *can be solved in polynomial time.*

Proof. We prove the theorem using the results of [9]. An alternative, stand-alone proof is given in Appendix A.

The input $\mathsf{PrXML}^{\mathsf{mux,ind,det}}$ document D can be rewritten to an equivalent $\mathsf{PrXML}^{\mathsf{exp}}$ document in polynomial time [2], which is a pTT document as defined by [9] (note that we make no use of the possibility of having uncertainty about order).

Furthermore, we can encode the deterministic document W to a deterministic tree automaton A_W with deterministic finite-state automata describing the regular languages of the transition function. Informally, the various states of the automaton will correspond to the various subtrees of W, except that subtrees occurring multiple times need to be identified. Formally, we define an equivalence relation $\sim$ on the nodes of W with $v \sim w$ if the subtrees rooted at v and w are isomorphic (i.e., they are the same tree, taking order into account. Let C_W be the set of classes of this relation, and ϕ be a mapping from the nodes of W to their class in C_W. We can use a dynamic bottom-up algorithm similar to that of the proof of Thm. 3 to compute the $\sim$ relation in polynomial time, as well as C_W and the mapping ϕ. Now, the alphabet of the automaton A_W is the set of node labels Λ, its set of states is C_W, its accepting state is $\phi(r)$ where r is the root of W, and its transition function maps $(c, l) \in C_W \times \Lambda$ to the empty language (if l is not the label of the nodes in c, noting that their labels must coincide) or (otherwise) to the language consisting of the single word $c_1 \cdots c_n$ where n is the number of children of all nodes v of W in the class c and, for all i, c_i is the class of the i-th child (note that n and the c_i do not depend on the choice of representative v). Computing A_W, with the languages of the transition function being represented by a deterministic finite-state automaton, can be done in polynomial time, and clearly by induction A_W accepts a tree T if and only if T is isomorphic to W.

The problem $\#\textsc{Poss}^{\mathsf{mux,ind,det}}_{<}$ then amounts to computing the total probability of the possible worlds of D that are accepted by A_W, which can be computed in polynomial time by Theorem 2 of [9]. □

4 Local models

We now complete the picture for the local model $\mathsf{PrXML}^{\mathsf{mux,ind,det}}$ on unordered documents. The results of [9] cannot be applied to this setting, as the ambiguity of node labels imply that we cannot impose an arbitrary order on document nodes; indeed, a reduction from perfect matching counting on bipartite graphs shows that the computation variant is hard even on the most inexpressive classes:

Theorem 2. $\#\textsc{Poss}^{\mathsf{ind}}_{\not\leq}$ *and* $\#\textsc{Poss}^{\mathsf{mux}}_{\not\leq}$ *are #P-hard, even when D has height* 4.

Proof. We first focus on the case of $\mathsf{PrXML}^{\mathsf{ind}}$. We show a reduction from the problem of counting the number of perfect matchings of a bipartite graph[4], which is #P-hard [21]. Let $G = (V, W, E)$ be a bipartite graph. We assume without loss of generality that $|V| = |W|$ (as G certainly cannot have perfect matchings otherwise), and let $n = |V| = |W|$.

Now, consider W with root labeled $\top$, n children labeled $\bot$, each of them with one child with labels respectively $l_1, \ldots, l_n$. Consider the uncertain document D with root labeled $\top$, n children labeled $\bot$, the i-th of them (for all i) having, for every j such that there is an edge in E from node i of V to node j of W, an ind child with one child labeled l_j with edge label $1/2$.

We claim that $D(W)$ is exactly the number of perfect matchings of the bipartite graph G, divided by $F = 2^{|E|}$.

To see why this is true, notice that each edge of E corresponds to an ind node of D. Hence, for any subset $M \subseteq E$, let us consider the valuation ν_M where the ind nodes for edges in M keep their child node, and the ind nodes for edges not in M discard their child node. This mapping between subsets of E and valuations is clearly one-to-one, and all those valuations have probability $1/F$ (because each of the $|E|$ events has probability $1/2$ and all of them are independent).

It only remains to see that the valuation ν_M yields W if and only if M is a perfect matching, but this is easy to see: if M is a perfect matching, each node labeled $\bot$ keeps exactly one child, and one node labeled l_i is kept for each node, so that ν_M yields W; conversely, if M is not a perfect matching, either there is a node labeled $\bot$ with zero or > 1 children, or there is some l_i kept zero or > 1 times, so that ν_M does not yield W. Hence, this completes the reduction, and shows that $\#\textsc{Poss}^{\mathsf{ind}}_{\not\leq}$ is #P-hard.

For the case of $\mathsf{PrXML}^{\mathsf{mux}}$, observe that the previous proof can be immediately adapted by replacing ind nodes with mux nodes, as every ind node has exactly one child. □

By contrast, the decision variant is tractable for $\mathsf{PrXML}^{\mathsf{ind}}$ and $\mathsf{PrXML}^{\mathsf{mux}}$, using a dynamic algorithm. However, allowing both ind and mux, or allowing det nodes, leads to intractability (by reductions from set cover and Boolean satisfiability).

Theorem 3. $\textsc{Poss}^{\mathsf{ind}}_{\top}$ *and* $\textsc{Poss}^{\mathsf{mux}}_{\top}$ *can be decided in PTIME.*

[4] Recall that a *perfect matching* in a bipartite graph is a subset of its edges such that each vertex of the graph (in either part) is adjacent to exactly one edge of the subset.

Proof. For ordered documents, the result follows from Theorem 1, so we only prove the claim that $\textsc{Poss}^{\mathsf{ind}}_{\not\leq}$ and $\textsc{Poss}^{\mathsf{mux}}_{\not\leq}$ can be decided in PTIME.

We show a stronger result, namely: the $\textsc{Poss}^{\mathsf{mux},\mathsf{ind}}_{\not\leq}$ problem can be decided in PTIME under the assumption that no ind node is a child of a mux node. Note that under this assumption, subtrees of D rooted at nodes that are not ind nodes only have possible worlds that are (possibly empty) subtrees (by contrast, ind nodes may have possible worlds that are forests). We say that a node of D is *non*-ind if it is a regular node or a mux node.

We will present a dynamic algorithm to decide $\textsc{Poss}^{\mathsf{mux},\mathsf{ind}}_{\not\leq}$ in PTIME under this assumption. We first compute bottom-up, for every non-ind node n of D, a Boolean value $e(n)$ indicating whether the subtree of D rooted at n has an empty possible world. If n is a regular node, we define $e(n) = \mathfrak{f}$. If n is a mux node, we define $e(n) = \mathfrak{t}$ if the probabilities of n sum up to < 1, or if one child n' of n is such that $e(n') = \mathfrak{t}$. It is clear that this computation can be performed in polynomial time.

The algorithm will now compute bottom-up, for every pair (n, n') of a non-ind node n in D and a node n' in W, a Boolean value $c(n, n')$ indicating whether or not the subtree of W rooted at n' is a possible world of the subtree of D rooted at n.

If n is a regular leaf, we define $c(n, n') = \mathfrak{t}$ if n' is a leaf with the same label as n, and $c(n, n') = \mathfrak{f}$ otherwise. Note that we can assume without loss of generality that all of D's leaves are regular nodes, as leaves that are probabilistic nodes can simply be removed.

If n is a mux node, we define $c(n, n') = \mathfrak{t}$ if one of the children x of n is such that $c(x, n')$ is $\mathfrak{t}$, otherwise $c(n, n') = \mathfrak{f}$. Observe that this is correct because the children of n are either mux nodes or regular nodes (they cannot be ind nodes), so the possible worlds of n are exactly the possible worlds of its children (possibly in addition to the empty subtree), and those possible worlds must be subtrees and not forests.

If n is an internal regular node of D, to define $c(n, n')$, we first check if n and n' have the same label. If they do not, we define $c(n, n') = \mathfrak{f}$; otherwise we continue.

Consider D the set of the topmost non-ind descendants of n. We say that a node x of D is *optional* if there is an ind node on the path from n to x, or if $e(x) = \mathfrak{t}$. In other words, a node x is optional if there is a valuation (of ind nodes) that discards it, or if there is a valuation of the subtree rooted at x which achieves an empty possible world for this subtree. This implies that, because the probabilistic choices are local and independent, we have a way to keep or delete every optional node of D *independently of each other*. Call D' the set of the children of n' in W.

Now if $|D| < |D'|$ we define $c(n, n') = \mathfrak{f}$ (because in no possible world can n have sufficiently many children to match n' – remember that the possible worlds of the subtrees of D rooted at non-ind nodes must be (possibly empty) subtrees but cannot be forests). Otherwise, add $|D| - |D'|$ *dummy nodes* to D' so that $|D'| = |D|$. Build a bipartite graph $G_{n,n'} = (D, D', E)$ with edges E defined as follows: an edge between $x \in D$ and non-dummy $x' \in D'$ if and only if $c(x, x')$ is 1, and an edge between x and dummy x' if and only if x was optional. (Intuitively: dummy nodes of D' represent the choice of deleting a node of D.)

We now claim that we should define $c(n, n') = \mathfrak{t}$ if and only if $G_{n,n'}$ has a perfect matching. To see why, observe first that $c(n, n')$ should be $\mathfrak{t}$ if and only if the subtree of W rooted at n' is a possible world of the subtree of D rooted at n, which, because n is

a regular node and the labels of n and n', amounts to saying that D' is a possible world of D. Observe now that for any subset S of E such that each vertex of D has exactly one incident edge, S describes a possible world of D: each node of D can achieve the node of D' to which it is thus matched (or, for dummy nodes, the empty subtree), because choices on the nodes of D (and their descendants, or at their parent edge in the case of deletions using an ind node) are independent between nodes of D. Now, a perfect matching describes a possible world of D achieving exactly D' (with no repetitions), and conversely if D' is a possible world of D it must be achieved by certain nodes of D realizing the nodes of D' (each node of D' being realized exactly once), and the others being deleted (each one being matched to one of the dummy nodes) so $G_{n,n'}$ must have a perfect matching.

Now, the existence of a perfect matching for the bipartite graph $G_{n,n'}$ can be decided in PTIME (using, e.g., the Hopcroft-Karp algorithm), so we can decide how to define $c(n,n')$ in PTIME (with a fixed polynomial not dependent on n or n').

Hence, we can decide in PTIME whether W is a possible world of D, by checking if $c(r,r')$ is $\mathfrak{t}$, with r and r' the roots of D and W (remember that r is assumed not to be a probabilistic node). This dynamic algorithm considers a quadratic number of pairs, and performs a polynomial-time computation (with a fixed polynomial) for each of them, so its overall running time is polynomial. □

Theorem 4. $\text{POSS}^{\mathsf{ind},\mathsf{det}}_{\not\leq}$, $\text{POSS}^{\mathsf{mux},\mathsf{det}}_{\not\leq}$ *and* $\text{POSS}^{\mathsf{mux},\mathsf{ind}}_{\not\leq}$ *are NP-complete, even when D has height* 4.

Proof. From Prop. 1 it suffices to show hardness. Let us first consider $\text{POSS}^{\mathsf{ind},\mathsf{det}}_{\not\leq}$. We show a reduction from the NP-hard [13] exact cover problem.

Consider an exact cover instance $S = \{S_1, \ldots, S_n\}$, where $S_i = \{s^i_1, \ldots, s^i_{n_i}\}$ for all i. Write $X = \bigcup S = \{v_1, \ldots, v_m\}$. The exact cover problem is to decide whether there exists a subset S' of S such that every element of X occurs in exactly one of the sets of S'.

Consider the document D with root labeled $\top$ and n ind children, with the i-th child having, for all i, only one child (with edge probability $1/2$), which is a det node, and which has n_i child nodes labeled $s^i_1, \ldots, s^i_{n_i}$. The document W has root labeled $\top$ and $|X|$ child nodes labeled $v_1, \ldots, v_m$.

W is a possible world if and only if there is some subset of the det nodes whose union yields exactly W (without duplicates), so that the reduction shows hardness.

To show hardness of $\text{POSS}^{\mathsf{mux},\mathsf{det}}_{\not\leq}$, observe that the previous proof can be adapted directly by replacing ind nodes by mux nodes, as every ind node has exactly one child.

Let us last consider $\text{POSS}^{\mathsf{mux},\mathsf{ind}}_{\not\leq}$. For this problem, we show a reduction from Boolean satisfiability. We use the same notations for the input instance as in the proof of Prop. 2. We additionally introduce n node labels $l_1, \ldots, l_n$, with label l_i corresponding to clause C_i.

Consider the document D whose root is labeled $\top$ and has m mux child nodes, each of them having two ind children with edge probability $1/2$, the probabilities of all edges of the ind nodes being also $1/2$. For all i, the first ind child of the i-th mux node has one child labeled l_j for every clause C_j where x_i occurs; the second one has one child labeled l_j for every clause C_j where $\neg x_i$ occurs. The document W has root labeled $\top$ and n children, the i-th one having label l_i.

We claim that W is a possible world of D if and only if $F = \bigwedge C_i$ is satisfiable. To see why, we consider a one-to-one mapping which associates, to any valuation ν of F, the outcomes of the mux nodes obtained by selecting the first child (resp. the second child) of the i-th mux node if $\nu(x_i) = \mathfrak{t}$ (resp. $\nu(x_i) = \mathfrak{f}$): by construction, the labels of the remaining ind nodes are those of the clauses which are true under valuation ν (possibly occurring multiple times). Hence, if there is a valuation ν satisfying F, then, selecting the outcomes of the mux nodes in this fashion, we can ensure that the remaining regular nodes are the $l_1, \ldots, l_n$, so that W is a possible world of D as we can choose a valuation of the ind nodes that keeps exactly one occurrence of each label.

Conversely, if W is a possible world of D, the outcome of the mux nodes in any outcome of D realizing W gives a valuation ν under which F is satisfied. Indeed, consider such an outcome and valuation ν, and, for any clause C_j of F, let us show that C_j is satisfied by ν. Because W is achieved, some node n labeled l_j must have been kept, and it must be the descendant of a mux node n' (say the i-th). Either it is a child of n''s first child n'_1, or of n''s second child n'_2. In the first case, this means that $\nu(x_i) = \mathfrak{t}$ because n'_1 was retained, and n being a child of n'_1 means that x_i occurs positively in C_j, so that C_j is true under ν. The second case is analogous. □

5 Explicit matches

We now attempt to understand how the overall hardness of POSS is caused by the difficulty of finding how the possible world W can be *matched* to D.

Definition 5. *A* candidate match *of W in D is an injective mapping f from the nodes of W to the regular nodes of D such that, if r is the root of W then $f(r)$ is the root of D, and if n is a child of n' in W then there is a descending path from $f(n)$ to $f(n')$ going only through probabilistic nodes.*

Intuitively, candidate matches are possible ways to generate W from D, ignoring probabilistic annotations, assuming we can keep exactly the regular nodes of D that are in the image of f. There are exponentially many candidate matches in general, so it is natural to ask whether POSS is tractable if all matches are explicitly provided as input:

Definition 6. *Given a class $\mathsf{PrXML}^{\mathcal{C}}$ and $o \in \{\bot, \not<, <, \top\}$, the* POSS *problem with* explicit matches $\mathrm{EPOSS}^{\mathcal{C}}_o$ *is the same as the $\mathrm{POSS}^{\mathcal{C}}_o$ problem except that the set of the candidate matches of W in D is provided as input (in addition to D and W).*

We study the explicit matches variant as a natural generalization of situations where the ways to match the possible world W to the document D are not too numerous and can be computed efficiently. For instance, if we assume that node labels in W are unique, so that there is no ambiguity about how to match W to D, then we are within the scope of the explicit matches variant, as the (unique) candidate match can be computed in polynomial time. The same applies to the situation where we only assume that no two sibling nodes carry the same label, or to more general settings where the possible matches can be identified easily. Requiring the possible matches to be provided as input is just a way to formalize that we are not accounting for the complexity of locating those matches.

We first note that explicit matches ensure tractability of all local dependency models, by reduction to deterministic tree automata [9], this time also for unordered documents. Intuitively, we can consider all candidate matches separately and compute the probability of each one, in which case no label ambiguity remains so any order can be imposed:

Theorem 5. $\#\text{EPOSS}^{\mathsf{mux,ind,det}}_{\top}$ *can be solved in polynomial time.*

Proof. We prove the theorem using the results of [9]. An alternative, stand-alone proof is given in Appendix B.

We say that a candidate match f is realized if we are in *the* possible world where the regular nodes of D that are kept are exactly those of the image of f. Hence, we can compute the probability of W by summing the probability of every candidate match being realized (because these events are mutually exclusive).

Now, to compute the probability of a candidate match f, replace the labels of nodes of W by unique labels (yielding W') and replace the labels of every node n of D in the image of f by the label of $f^{-1}(n)$ in W', to obtain a probabilistic document D'. The probability of f being realized is $D'(W')$. Importantly, if D and W are unordered, we can make D' and W' ordered by choosing any order on sibling nodes in D', and apply the same order (following f^{-1}) to sibling nodes in W; this works because the way to match W to D is fully specified by f so there is no matching ambiguity when imposing this order.

This concludes the proof, because D' and W' are computable in polynomial time and $D'(W')$ can be computed by a deterministic tree automaton as in the proof of Theorem 1. □

For long-distance dependencies, however, it is easily seen that POSS is still hard with conjunction of events, even if explicit matches are provided:

Theorem 6. $\text{EPOSS}^{\mathsf{cie}}_{\bot}$ *is NP-complete, even when D has height* 3.

Proof. From the proof of Prop. 2, noticing that there is only one (trivial) match of W in D for the instances considered in the reduction. □

This being said, it turns out that the hardness is really caused by event *conjunctions*. To see this, we introduce the $\mathsf{PrXML}^{\mathsf{mie}}$ class, which allows only *individual* events:

Definition 7. *The* $\mathsf{PrXML}^{\mathsf{mie}}$ *class features* multivalued independent events *taking their values from a finite set V (beyond* ⊺ *and* ⊺*, with probabilities summing to* 1*), and probabilistic* mie *nodes whose child edges are annotated by a* single *event e and a value* $x \in V$*. A* mie *node cannot be the child of a* mie *node. When evaluating D under a valuation* ν*, child edges of* mie *nodes labeled* (e,x) *should be kept if and only if* $\nu(e) = x$*.*

Note that mie hierarchies are forbidden (because they can straightforwardly encode conjunctions), so that $\mathsf{PrXML}^{\mathsf{mie}}$ does not capture ind hierarchies. However, as we introduced it with multivalued (not just Boolean) events, it captures $\mathsf{PrXML}^{\mathsf{mux}}$:

Proposition 3. *We can rewrite* $\mathsf{PrXML}^{\mathsf{mux}}$ *to* $\mathsf{PrXML}^{\mathsf{mie}}$ *and* $\mathsf{PrXML}^{\mathsf{mie}}$ *to* $\mathsf{PrXML}^{\mathsf{cie}}$ *in PTIME.*

Proof. To rewrite $\mathsf{PrXML}^{\mathsf{mux}}$ to $\mathsf{PrXML}^{\mathsf{mie}}$, first rewrite the input $\mathsf{PrXML}^{\mathsf{mux}}$ document to a $\mathsf{PrXML}^{\mathsf{mux}}$ document with no mux hierarchies (no mux node is a child of a mux node); this can be done in polynomial time ([2], Lemma 5.1). Next, introduce one event per mux node and one outcome for this event per child of the mux, with one additional outcome (to make the probabilities sum to 1) if the original probabilities of the mux child edges summed to < 1. Replace each mux node by a mie node, where every child edge of the mie node is labeled by the event introduced for this mux node and the value introduced for the outcome where this child edge is kept. The absence of mux hierarchies ensures that the requirement on the absence of mie hierarchies is respected.

To rewrite $\mathsf{PrXML}^{\mathsf{mie}}$ to $\mathsf{PrXML}^{\mathsf{cie}}$, we claim that every multivalued event e with k outcomes can be replaced by a set S_e of $O(k)$ Boolean events such that each outcome $e = x_i$ can be represented by a conjunction of $O(\log_2 k)$ events of S_e, those conjunctions having the same probability as their original outcome and forming a partition of all outcomes of events in S_e. Assuming that this claim holds, the $\mathsf{PrXML}^{\mathsf{mie}}$ document can be rewritten in polynomial time to $\mathsf{PrXML}^{\mathsf{cie}}$ by performing this encoding for all multivalued events, and replacing every mie node by a cie node and replacing each child edge labeled (e, x_i) by a child edge labeled with the corresponding conjunction.

Now, to see why the claim is true, given a multivalued event e, observe that we can build a binary decision tree T_e of the outcomes of e. Hence, we can introduce one Boolean event per internal node of T_e, and choose its probability according to that of its two child edges in T_e (the probability of an edge a in T_e being the total probability of the outcomes reachable from the target of a, normalized by that of the outcomes reachable from the origin of a). Hence, we associate to each outcome x_i of e the conjunction of Boolean choices leading to x_i in T_e: it has the right probability by construction, and, for every valuation of the Boolean events, exactly one conjunction is true (the one corresponding to the leaf of T_e selected by following those choices). Now, as T_e is a binary tree with k leaves (the number of outcomes of e), it has $O(k)$ internal nodes and its height is $O(\log_2 k)$, which proves the claim and completes the proof.

Observe that $\mathsf{PrXML}^{\mathsf{mie}}$ does *not* capture $\mathsf{PrXML}^{\mathsf{mux,det}}$; a proof of this fact is given in Appendix C. □

In the $\mathsf{PrXML}^{\mathsf{mie}}$ class, the POSS problem is still NP-hard, by reduction to exact cover; however, with explicit matches, the #POSS problem is tractable, both in the ordered and unordered setting, despite the long-distance dependencies. Intuitively, the candidate matches are mutually exclusive, and each match's probability can be computed as that of a conjunction of equalities and inequalities on the events at the frontier.

Theorem 7. $\text{POSS}_{\perp}^{\mathsf{mie}}$ *is NP-complete, even when D has height* 3 *and events are Boolean.*

Proof. From Prop. 1 it suffices to show hardness. We show a reduction from exact cover, as in the proof of Theorem 4, with the same notation for the exact cover instance (and, intuitively, using for D and W the straightforward encoding to $\mathsf{PrXML}^{\mathsf{mie}}$ of the instances used in this last proof to show hardness of $\text{POSS}_{\perp}^{\mathsf{mux,det}}$ and $\text{POSS}_{\perp}^{\mathsf{ind,det}}$).

Consider a set of n Boolean events $E = \{e_1, \ldots, e_n\}$ (with values in $\{\mathrm{t}, \mathrm{f}\}$ and probabilities $1/2$. Consider the document W with one root labeled $\top$ and m children labeled $l_1, \ldots, l_m$. Consider the $\mathsf{PrXML}^{\mathsf{mie}}$ document D with one root labeled $\top$ and one mie child with, for $1 \leq j \leq n$, n_i child edges labeled (e_i, t) leading to children labeled

$s^i_1, \ldots, s^i_{n_i}$. Order in the input D the child nodes of the root in W from l_1 to l_m, and the child nodes of the root in D from those labeled l_1 to those labeled l_m, the order between those carrying the same labels being arbitrary, so that we are showing a reduction either to $\text{POSS}^{\text{mie}}_{\not<}$ or to $\text{POSS}^{\text{mie}}_{<}$.

Now, W is a possible world of D if and only if there is a valuation of the events of E such that, for every $1 \leq j \leq m$, there is exactly one node labeled l_j that is retained. This amounts to choosing a subset S' of S such that every item of X occurs exactly once in $\bigcup S'$: the set S' corresponds to the set of events of E that are evaluated to $\mathfrak{t}$. Hence, (D, W) is a positive instance of $\text{POSS}^{\text{mie}}_{\not<}$ if and only if F is satisfiable, so that $\text{POSS}^{\text{mie}}_{\perp}$ is NP-hard. □

Theorem 8. $\#\text{EPOSS}^{\text{mie}}_{\top}$ *can be solved in polynomial time.*

Proof. Observe first that, as in the proof of Theorem 5, the probability that W is realized is that of either of the candidate matches being realized, those events being mutually exclusive. We assume that, if W and D are ordered, we have checked (in PTIME) that candidate matches respect the order (for a candidate match f, if v and v' are sibling nodes in W such that v comes before v', then $f(v)$ comes before $f(v')$ in the document order of D), and removed those which do not.

Now, consider a candidate match f. We must compute the probability p_f that f is realised, namely, that we are in the possible world where the only regular nodes that are kept in D are those of the image I of f; we abuse notation so that we consider mie nodes of D to be in I if one of their children is in I. We will write this probability p_f as that of a conjunction of events: the events that all nodes in I are kept, and the events that all nodes not in I are discarded.

The event of all nodes in I being kept can be written as the conjunction c_+ of all $e_i = x_i$ for every edge (e_i, x_i) between a mie node in I and a child node also in I. Indeed, to keep I, all the conditions on edges leading to a node of I must be respected.

The event of all nodes not in I being discarded can be written as a conjunction c_- of the same kind, in the following fashion. Consider every topmost node n not in I. If n's parent n' is a regular node, then the overall probability of the match f is $p = 0$, because if we keep n' then we must keep n; in this case, we can forget about f altogether. Otherwise, we add to c_- the atom $e_i \neq x_i$, where (e_i, x_i) is the label of the edge from n' to n.

We now have either eliminated f or obtained (in polynomial time) the conjunction $c = c_+ \wedge c_-$ which is necessary and sufficient for the match to hold, the atoms of c being of the form $e_i = x_i$ or $e_i \neq x_i$, where the e_i's are events and the x_i's are outcomes. Now, we can compute in polynomial time the probability p_f of c. Indeed, regroup the atoms by the probabilistic event occurring in them. For each probabilistic event e, we consider the (possibly empty) subset of outcomes satisfying the atoms for e, and compute its total probability p^e_f. As the choices are independent between events, the overall probability p_f of c is the product of the p^e_f over all events e. □

6 Conclusion

We have characterized the complexity of the counting and decision variants of POSS for unordered or ordered XML documents, and various PrXML classes. With explicit

matches, #POSS is tractable unless event conjunctions are allowed. Without explicit matches, POSS is hard unless dependencies are local; in this case, if the documents are ordered, #POSS is tractable, otherwise #POSS is hard and POSS is tractable only with ind *or* mux nodes (and hard if both types, or det nodes, are allowed). Our results are summarized in Table 1 on page 5.

We note that, using our results and via translations between the probabilistic relational and XML models [4], we can derive some bounds on the complexity of POSS for relational databases. In terms of tractability for the (unordered) relational model, we can deduce the tractability of the decision formulation of POSS for the tuple-independent model [17,10] and the block-independent-disjoint model [7,19], and the tractability of both the decision and counting variants on pc-tables [11,12] under the assumption that explicit matches are provided and that tuples are annotated by a single equality constraint on a multivalued event, in the spirit of mie. We remark, however, that such results are not hard to prove directly in the relational model. In terms of intractability, we observe that the translation from XML to relational models in [4] requires the introduction of explicit node IDs for all nodes of the document, so that this does not translate to a reduction for the POSS problem: intuitively, the translation of W to a relational table would have to specify the exact node IDs to be matched. We leave as future work a more complete investigation of POSS in the relational context, or the study of possible alternative translations that provide more reductions for POSS from one setting to the other.

Additional directions for future work would be to study more precisely the effect of det nodes and ind hierarchies, for instance by attempting to extend the $\mathsf{PrXML}^{\mathsf{mie}}$ class to capture them, or try to understand whether there is a connection between the algorithms of [9] and the proof of Thm. 3. It would also be interesting to determine under which conditions (beyond unique labels) can candidate matches be enumerated in polynomial time, so that the POSS problem reduces to the explicit matches variant. Last but not least, another natural problem setting is to allow the order on sibling nodes of D to be partly specified. This question is already covered in [9], but only when all of the possible orderings are explicitly enumerated: investigating the tractability of POSS for more compact representations, such as partial orders, is an intriguing problem.

Acknowledgements. The author thanks Pierre Senellart for careful proofreading, useful suggestions, and insightful feedback, the anonymous referees of AMW 2014 and BDA 2014 for their valuable comments, and M. Lamine Ba and Tang Ruiming for helpful early discussion. This work has been partly funded by the French government under the X-Data project and by the French ANR under the NormAtis project.

References

1. S. Abiteboul, T.-H. H. Chan, E. Kharlamov, W. Nutt, and P. Senellart. Capturing continuous data and answering aggregate queries in probabilistic XML. *ACM Transactions on Database Systems*, 36(4), 2011.
2. S. Abiteboul, B. Kimelfeld, Y. Sagiv, and P. Senellart. On the expressiveness of probabilistic XML models. *VLDB Journal*, 18(5):1041–1064, 2009.
3. A. Amarilli. The possibility problem for probabilistic XML. In *Proc. AMW*, 2014.

4. A. Amarilli and P. Senellart. On the connections between relational and XML probabilistic data models. In *Proc. BNCOD*, pages 121–134, Oxford, United Kingdom, 2013.
5. M. L. Ba, T. Abdessalem, and P. Senellart. Merging uncertain multi-version XML documents. *Proc. DChanges*, 2013.
6. M. L. Ba, T. Abdessalem, and P. Senellart. Uncertain version control in open collaborative editing of tree-structured documents. In *Proc. DocEng*, pages 27–36, 2013.
7. D. Barbará, H. Garcia-Molina, and D. Porter. The management of probabilistic data. *IEEE Transactions on Knowledge and Data Engineering*, 4(5), 1992.
8. P. Barceló, L. Libkin, A. Poggi, and C. Sirangelo. XML with incomplete information. *JACM*, 58(1):4, 2010.
9. S. Cohen, B. Kimelfeld, and Y. Sagiv. Running tree automata on probabilistic XML. In *Proc. PODS*, pages 227–236. ACM, 2009.
10. N. N. Dalvi and D. Suciu. Efficient query evaluation on probabilistic databases. *VLDB Journal*, 16(4), 2007.
11. T. J. Green and V. Tannen. Models for incomplete and probabilistic information. In *Proc. EDBT Workshops, IIDB*, Mar. 2006.
12. J. Huang, L. Antova, C. Koch, and D. Olteanu. MayBMS: a probabilistic database management system. In *SIGMOD*, 2009.
13. R. M. Karp. *Reducibility among combinatorial problems*. Springer, 1972.
14. E. Kharlamov, W. Nutt, and P. Senellart. Updating probabilistic XML. In *Proc. Updates in XML*, Lausanne, Switzerland, 2010.
15. B. Kimelfeld, Y. Kosharovsky, and Y. Sagiv. Query evaluation over probabilistic XML. *VLDB Journal*, 18(5):1117–1140, 2009.
16. B. Kimelfeld and P. Senellart. Probabilistic XML: Models and complexity. In Z. Ma and L. Yan, editors, *Advances in Probabilistic Databases for Uncertain Information Management*, pages 39–66. Springer-Verlag, 2013.
17. L. V. S. Lakshmanan, N. Leone, R. B. Ross, and V. S. Subrahmanian. ProbView: A flexible probabilistic database system. *TODS*, 22(3), 1997.
18. L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
19. C. Ré and D. Suciu. Materialized views in probabilistic databases: for information exchange and query optimization. In *VLDB*, 2007.
20. D. Suciu, D. Olteanu, C. Ré, and C. Koch. *Probabilistic Databases*. Morgan & Claypool, 2011.
21. L. G. Valiant. The complexity of computing the permanent. *Theoretical computer science*, 8(2):189–201, 1979.

A Stand-alone proof of Theorem 1

We prove the claim by representing ordered trees as words (essentially following a SAX traversal). First, encode W to a word e_W by such a traversal, internal nodes with label a being encoded as $a^l C a^r$ where C is the sequence of the encodings of the children of the node (following their order), and leaves with label a being encoded as $a^l a^r$.

We now convert D to a weighted non-deterministic automaton A_D (on words) with ε-transitions; importantly, this automaton is acyclic. We proceed in the following way. Encode a regular node n with label a as the following structure: the initial state q_i, the encoding of the children (c_i) of n in order (the final state of each one being connected to the initial state of the next one by an ε-transition of weight 1), the final state q_f, and an edge labeled a^l with probability 1 from q_i to the initial state of the encoding of c_1 (if it

exists, otherwise to some intermediate state q) and an edge labeled a^r with probability 1 from the final state of the encoding of the last child (if it exists, otherwise from q) to q_f.

Encode the det nodes in the same way except that the two last edges are labeled by ε (instead of a^l and a^r). Encode an ind node n like a regular node except that edges leading to the initial state of the encoding of a child of n are given a probability p (the probability of this child) and we add an additional edge with label ε and probability $1-p$ to the same initial state to the final state of the encoding of that child (corresponding to the choice of not retaining this child). Encode a mux node as an initial state q_i, an initial state q_f, the encoding of each child in parallel, ε-transitions with probability 1 from the final state of the encoding of the children to q_f, and ε-transitions with adequate probabilities from q_i to the initial state of the encoding of each child (or to q_f, to make the probabilities sum to 1).

There is a clear correspondence between runs of A_D and possible worlds of D, so that what we have to compute is the probability of the encoding e_W of W according to A_D.

Now, because A_D is acyclic, it is easy to compute this probability in polynomial time. Indeed, we can compute dynamically for each state q of A_D and every suffix s of e_W the probability that s is produced by a run from q to the final state of A_D.

The base case is that, at the final state, we produce the empty suffix with probability 1 and any non-empty suffix with probability 0.

Now, when considering a non-final state q and suffix s, because by construction the sum of all outgoing transitions of q is 1, the probability $p(q,s)$ of producing s from q is computed by summing, for every outgoing transition a starting at state q (with target state q_i), the probability of the transition a multiplied by the following quantity: either, if a is an ε-transition, the value $p(q_i,s)$ (which was already computed) or, if a has label x, either 0 if $|s|=0$ or the first letter of s is not x, or otherwise the value $p(q_i,s')$ (which was already computed) where s' is the suffix of length $|s|-1$ of e_W.

B Stand-alone proof of Theorem 5

As in the other proof of this result, it suffices to consider a single candidate match, as the overall probability can be obtained by summing that of each match, and the decision problem can be solved by considering matches separately. If D and W are ordered, we can assume that matches which do not satisfy the order constraints have been discarded. For simplicity we relax the restriction that the probability of edges is always < 1, so that we can encode det nodes as ind nodes and consider only ind and mux nodes.

Let us first prove that $\text{EPOSS}_{\top}^{\mathsf{mux,ind,det}}$ can be solved in polynomial time.

Consider a candidate match f. Because all probabilistic choices are independent, it is clear that all nodes of D in the image I of f can be kept if and only if there is no mux node n such that n' and n'' are in I and n', n'' are descendants of two distinct children of n. Indeed, this condition is clearly necessary, and, except for this, all choices are independent within I so they can all be made to succeed[5] so that the nodes of I are retained.

[5] Remember that there are no edges with probability 0.

So, assuming that this condition holds (it can be checked in polynomial time), the question is only to see whether the nodes not in I can be discarded. To check this, we define recursively on all nodes n of D, in a bottom-up fashion, the Boolean value $e(n)$ indicating if n can be "empty", that is, if there is a possible world rooted at n that is empty.

If n is a regular node then we define $e(n) = \mathfrak{f}$.

If n is an ind node, we define $e(n) = \mathfrak{t}$ if and only if $e(n)$ is $\mathfrak{t}$ for all the children of n with edge probability 1 (remember that we relaxed the condition on probabilities being < 1 because we encoded det nodes as ind nodes). In particular, if n has no children with edge probability 1, we define $e(n) = \mathfrak{t}$.

If n is a mux node, we define $e(n) = \mathfrak{t}$ if and only if the probabilities of n sum up to < 1, or there exists a child of n such that $e(n)$ is $\mathfrak{t}$.

Now, we can use e to express the fact that it should be possible to discard all regular nodes of D except those in I. To do so, by a slight abuse of terms, we say that a probabilistic node is in I if it has a regular descendant that is in I. Now, we claim that the match can yield W if and only if, for every topmost node n not in I, either $e(n)$ is true, or the parent of n (which by definition is in I) is a mux node or is an ind node n' such that the edge from n' to n is labeled with a probability < 1. To see why this claim holds, observe that, if this condition is respected, all nodes not in I can be discarded (either because $e(n)$ is $\mathfrak{t}$ so we can choose an empty subtree as the possible world rooted at them, or by deciding to discard them at the level of their parent – for mux nodes, in fact, we have no choice but to discard them). Conversely, if this condition does not hold for a node n' of D, then n' must be kept, and the possible world chosen for the subtree rooted at n' will have to be non-empty because $e(n')$ is $\mathfrak{f}$.

This condition can be tested in polynomial time. Hence, $\text{EPOSS}_{\top}^{\mathsf{mux,ind,det}}$ can be solved in polynomial time by checking if one of the matches is acceptable in this sense.

Let us now prove that $\#\text{EPOSS}_{\top}^{\mathsf{mux,ind,det}}$ can be solved in polynomial time. We assume that candidate matches are filtered (in polynomial time) according to the process described above, so as to only keep the matches with probability > 0.

Now, the probability that the match f is realized can be computed as the probability of keeping its image I (including probabilistic nodes like in the previous proof), times the probability of discarding the other nodes: indeed, as I is a rooted subtree, we must first decide outcomes of nodes and edges in this subtree so that I is kept, and then outcomes such that the rest is discarded.

It is easily seen that the probability p_+ that I is kept is the product of all probabilities that annotate the edges that are between nodes in I: all ind edges of the match must be kept (and those draws are performed independently), and the right mux edges must always have been chosen (remember that a mux node n is in I only if it has a descendant in I, and by the condition that the match probability is > 0 all descendants of n are descendants of the same child of n).

Now, we must compute the probability p_- that the nodes not in I are discarded. To do so, we define $e(\cdot)$, as in the previous proof, but as a probability rather than a Boolean value: $e(n)$ is the probability of the empty subtree among the possible worlds for the subtrees rooted at $e(n)$ (note that this probability does not depend on the choices performed elsewhere in the tree). Once again, we compute $e(\cdot)$ bottom-up.

For a regular node n, we define $e(n) = 0$.

For a mux node n, we define $e(n) = (\sum_i p_i e(n_i)) + (1 - \sum_i p_i)$ where the n_i are the children of n and the p_i the corresponding edge labels. Intuitively, the probability of the mux to be empty is that of its children being empty, weighted by their probability, plus the probability that we select no children (when the probabilities sum to < 1).

For an ind node n, we define $e(n) = \prod_i ((1 - p_i) + p_i e(n_i))$ with the same notation. Intuitively, the probability of the ind to be empty is that of each child subtree being missing or empty, which occurs either when the corresponding is removed, or when it is kept but the subtree is empty (summing those two cases are they are mutually exclusive).

Now, all nodes not in I are discarded if and only if, for each topmost node n not in I, either n is dropped (its parent edge is removed) or the possible world rooted at n is empty. This is a conjunction of events, and they are independent once conditioned by the fact that the nodes in I are kept (so the outcomes of all mux nodes in I have already been decided), so we can compute the overall probability of f as $p_+ p_-$, with p_- being the product of the probability p_n, for all topmost nodes n not in I, that n is dropped or the subtree rooted at n is empty. Consider n' the parent of n: the probability p_n is $e(n)$ if n' is a regular node (as n cannot be dropped then), is 1 if n' is a mux node (as n' is in I, it has a descendant n'' in I, which cannot be a descendant of n as n is not in I, so that when deciding to keep n'' we have already decided that n would be dropped), and it is $pe(n) + (1 - p)$ if n' is an ind node and the probability of the edge from n' to n is p.

Hence, the overall probability, $p_+ p_-$, can be computed in polynomial time, which concludes the proof.

C $\mathsf{PrXML}^{\mathsf{mie}}$ does not capture $\mathsf{PrXML}^{\mathsf{mux,det}}$

In this section, we show that $\mathsf{PrXML}^{\mathsf{mie}}$ is *not* more general than $\mathsf{PrXML}^{\mathsf{mux,det}}$, namely, there is no PTIME encoding from $\mathsf{PrXML}^{\mathsf{mux,det}}$ documents to $\mathsf{PrXML}^{\mathsf{mie}}$ documents.

Consider the $\mathsf{PrXML}^{\mathsf{mux,det}}$ document D_n with root labeled $\top$ and one mux child that has two children (with edge probabilities $1/2$): one regular child with label c, and one det child. The det node has n mux children: for all i, the i-th of them has edge probabilities $1/2$ and two regular children with labels a_i and b_i. We show that any encoding of D_n to a $\mathsf{PrXML}^{\mathsf{mie}}$ document D'_n (having the same possible worlds as D_n) must have size exponential in n.

The document D'_n must have root labeled $\top$, and the root clearly cannot have any regular children; so it must have mie children, and without loss of generality it has only one of them. Now, as mie hierarchies are not permitted, all children of this mie node are regular nodes; clearly they cannot have any regular children, and without loss of generality they have no (useless) mie children. So the only thing to define is the label and edge labels of the children of this unique mie node. Without loss of generality we assume that we remove edges labeled with (e, v) where $e = v$ has probability 0. Clearly the node labels can be assumed to be either c or a_i or b_i for some i.

As all possible worlds of D'_n must contain at most one node labeled c, we claim that the parent edge of all child nodes with label c must be labeled with the same event e. Indeed, if there are two nodes with label c and with edge labels (e_1, v_1) and (e_2, v_2)

with $e_1 \neq e_2$, because e_1 and e_2 are independent and (we assumed) the events $e_1 = v_1$ and $e_2 = v_2$ have probability > 0, any valuation where $e_1 = v_1$ and $e_2 = v_2$ yields a possible world with two c children, a contradiction. Hence all child nodes with label c are labeled with the same event e. Note that, as some possible world of D'_n must contain a node labeled c, there has to be at least one child n_c with label c.

Now, no possible world of D'_n contains both a child labeled a_i or b_i (for any i) and a child labeled c, so we claim that the parent edge of all child nodes with label a_i or b_i (for any i) must be labeled with this same event e. Indeed, assume that one such node is labeled with some condition (e', x) with $e' \neq e$; calling (e, v) the edge label of n_c, any valuation where $e' = x$ and $e = v$ yields a possible world with a c node and an a_i node or a b_i node for some i, a contradiction. Hence, in fact, all child nodes of the mie node are labeled with the same event e.

Now, as D_n has $2^n + 1$ possible worlds, e must have $\Omega(2^n)$ different possible values. Hence, D'_n is of size exponential in n. This concludes the proof.